Available online at www.sciencedirect.com
ScienceDirect
Journal of Non-Crystalline Solids 352 (2006) 4424–4429
JOURNAL OF NON-CRYSTALLINE SOLIDS
www.elsevier.com/locate/jnoncrysol

# Exploiting memory in event-based simulations

Mykyta V. Chubynsky [a], H. Vocks [b], G.T. Barkema [b], Normand Mousseau [a,b,*]

[a] *Département de Physique and Regroupement québécois sur les matériaux de pointe, Université de Montréal, C.P. 6128, Succursale Centre-ville Montréal, Québec, Canada H3C 3J7*
[b] *Institute for Theoretical Physics, Utrecht University, Leuvenlaan 4, 3584 CE Utrecht, The Netherlands*

Available online 17 August 2006

**Abstract**

Few simulation methods have succeeded in sampling efficiently the phase space of complex systems with a dynamics dominated by activated events. In order to address this limitation, we have recently introduced an activated algorithm based on a mixture of the activation-relaxation technique and molecular dynamics (the properly obeying probability activation relaxation technique, POP-ART). We show here that the basic implementation of POP-ART is only as fast as MD in sampling the phase space of a complex material, amorphous silicon at 600 K. However, as the activation moves are locally defined, it is possible to use a number of tricks that can increase significantly sampling efficiency of POP-ART. We consider an approach, the memory kernel, based on avoiding recently encountered moves and show using a simple model that this introduces very little bias while ensuring a significant gain over standard Monte Carlo in sampling the phase space of this model. Incorporating the memory kernel into POP-ART improves considerably its efficiency in sampling the phase space of amorphous silicon as compared to standard POP-ART and molecular dynamics.


*PACS:* 02.70.−c; 61.43.−j; 61.43.Bn; 82.20.Wt

*Keywords:* Amorphous semiconductors; Glasses; Modeling and simulation; Molecular dynamics; Monte Carlo simulations

## 1. Introduction

Considerable effort has been invested over the last decade to develop algorithms that can sample the phase space of complex materials with a dynamics dominated by rare activated events. Thermodynamical methods such as simulated tempering [1] and entropic sampling and related methods [2,3] bias the sampling by exchanging configurations at different temperatures or sampling states based on a preselected distribution. In parallel, a number of dynamical methods have also been proposed, such as hyper-MD [4] and self-guided dynamics [5], to be applied to systems dominated by rare events.

Another class of algorithms is based on identifying pathways leading from one local minimum to another [6–10]; the current algorithms in this class do not guarantee proper thermodynamical sampling. Recently, we have introduced a new method in this class, the properly obeying probability activation-relaxation technique (POP-ART) [11] which likewise generates activated trajectories through the configuration space, but now ensuring correct sampling. Based on a mixture of molecular dynamics and Monte Carlo, this method can integrate a number of MC *tricks* which can improve significantly its efficiency.

In this paper, we revisit POP-ART, showing how the sampling efficiency can be improved through the use of a memory kernel. We study first the impact of this memory kernel in a simple toy model, and then in a simulation of the relaxation of a 1000-atom model of amorphous silicon at 600 K. We stress that the memory kernel approach is not limited to POP-ART and could be a useful addition to many accelerated methods already in place.

* Corresponding author.
*E-mail addresses:* mykyta.chubynsky@umontreal.ca (M.V. Chubynsky), vocks@phys.uu.nl (H. Vocks), barkema@phys.uu.nl (G.T. Barkema), normand.mousseau@umontreal.ca (N. Mousseau).

0022-3093/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.jnoncrysol.2006.01.117

## 2. The properly obeying probability activation-relaxation technique

Before describing the improvements to POP-ART, we summarize the original algorithm here. A more extensive discussion of the method as well as a full justification of the various steps can be found in Ref. [11].

POP-ART divides the system's configuration space into two types of regions: the basins and the saddle region. The basins surround local energy minima while the saddle region includes first and higher-order saddle points. The threshold separating the basins from the saddle region is based on the value of the second derivative of the potential energy in the direction in which it is lowest, denoted $\lambda$. This value is obtained by diagonalizing the Hessian. The critical threshold $\lambda_0$ is chosen such that the system spends most of its time in the basin region and only rarely crosses into the saddle region. In general, therefore, $\lambda_0$ is negative.

Using this difference between regions, POP-ART defines two types of motion. Motion in the basins is performed using conventional molecular dynamics. Once the system reaches the boundary, however, the MD is stopped and the system is brought to a new point on a basin boundary going through the saddle region following rules defined below. Once at the boundary of a new basin, the molecular dynamics is resumed. All steps in POP-ART are designed to ensure detailed balance and, therefore, proper thermodynamical sampling.

The originality of POP-ART lies in the activated move. It is composed of two steps: (1) a trajectory is first generated, connecting two points on the basin boundaries; (2) the free energy difference between these two points is then evaluated and the move is accepted or rejected according to the Metropolis criterion. Without going into the details, we discuss below these two steps.

The first step, generating the activated trajectory, is performed in the spirit of the activation-relaxation technique (ART) [6,10]. More precisely, the trajectory follows the direction of the eigenvector associated with the lowest eigenvalue of the Hessian and pointing away from the basin. However, because we intend to simulate the microcanonical ensemble, the move is defined *at constant energy*, i.e., the activated trajectory is generated by a series of discrete steps according to

$$\vec{x}_{i+1} = \vec{x}_i + \frac{\Delta\tau}{2}(\vec{h}_i + \vec{h}_{i+1}) + \frac{c\Delta\tau}{2}(\vec{F}_i + \vec{F}_{i+1}), \tag{1}$$

where $\vec{h}_i$ is the normalized eigenvector at $\vec{x}_i$ corresponding to the lowest Hessian eigenvalue, $\vec{F}_i$ is the total $N$-dimensional force at $\vec{x}_i$, $\Delta\tau$ is a constant step, and $c$ is a multiplicative constant, chosen to ensure constant potential energy. The values of $\vec{h}$ and $\vec{F}$ at $\vec{x}_{i+1}$ are obtained iteratively. The activation process is continued until the eigenvalue associated with $\vec{h}_i$ reaches the threshold $\lambda_0$. This equation is valid when all atomic masses are equal; it is straighforward to extend it to a system with multiple atomic species.

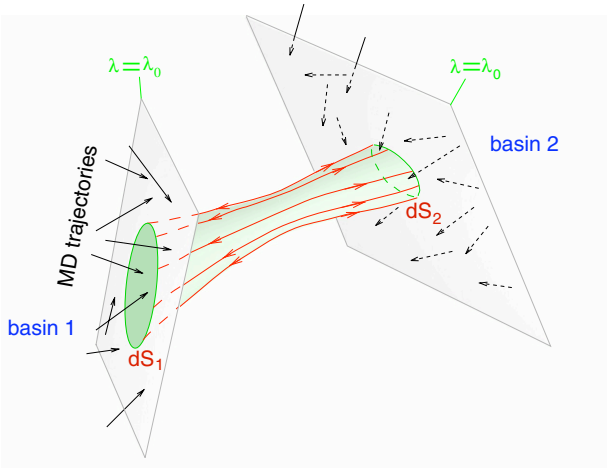


Fig. 1. Sketch of tube connecting two basins. The planes are at $\lambda = \lambda_0$ and the tube represents a family of trajectories obtained by moving slightly the entrance point. The width of this tube is associated with the cross-section Jacobian. The areas of the sections of the tube on the constant eigenvalue surfaces $\lambda = \lambda_0$ are denoted by $dS_1$ and $dS_2$.

While the trajectory defined here is symmetric and keeps the total energy constant, we must nevertheless compute the change in free energy between the two endpoints. We compute the Jacobian $J$ of the transformation associated with the chosen trajectory. This quantity represents the expansion or the contraction of the volume formed by nearby points in configuration space as the system is pushed along the activated trajectory. Fig. 1 sketches a trajectory and the main contributions to the free energy difference: the boundary and the cross-section contributions. The first contribution is the boundary Jacobian, $J_b$, which is due to the fact that the activation trajectory does not start and end perpendicular to the boundary but at some angles $\alpha_1$ and $\alpha_2$, respectively. For this reason, the ratio of the area of the exit from the tube to that of the entrance differs from the ratio of the cross-sections of the tube at the exit and at the entrance by a factor [11]

$$J_b = \frac{\cos\alpha_1}{\cos\alpha_2} \tag{2}$$

with

$$\cos\alpha_{1,2} = \frac{\vec{h}\cdot\nabla\lambda}{|\nabla\lambda|}, \tag{3}$$

and

$$\nabla\lambda(\vec{x}) = \lim_{\delta\to 0}\frac{2\vec{F}(\vec{x}) - \vec{F}(\vec{x}+\delta\vec{h}) - \vec{F}(\vec{x}-\delta\vec{h})}{\delta^2}, \tag{4}$$

where $\lambda$ is the lowest eigenvalue of the Hessian and its gradient is taken at the entrance and exit points for $\alpha_1$ and $\alpha_2$, respectively.

Similarly, the logarithm of the cross-section Jacobian $J_{xs}$ is an integral along the trajectory:

$$\ln J_{xs} = \int_0^\tau j(\vec{x}(\tau'))d\tau', \tag{5}$$

where $j$ is given by

$$j = \operatorname{div}\vec{h} + c\operatorname{div}\vec{F}. \tag{6}$$

Therefore a POP-ART run follows these steps:

(1) We start in a basin doing MD at the constant energy of interest. Every few steps, we compute the lowest eigenvalue of the Hessian using the Lanczos algorithm.
(2) When the boundary is reached, the MD is stopped, the velocity stored and the activated trajectory is generated using Eq. (1) with the step $\Delta\tau = 0.01$ Å.
(3) We compute the boundary factor and the cross-section Jacobian and accept the event (the move from the beginning of the activation trajectory to its end) with a probability $\min(1, J)$.
(4) If the event is accepted, we start from the endpoint of the activated trajectory using the initial velocity (with the component normal to the boundary reversed, if necessary, to ensure that the velocity points inside the basin). Otherwise, we restart from the point of entry, reversing the normal component of the velocity with respect to the boundary.
(5) Once the threshold is reached again, we repeat steps 2–5.

## 3. The memory kernel

POP-ART, as described above, has been shown to sample correctly the thermodynamical states associated with an interstitial in Si [11]. Moreover, in simple systems, with a well defined barrier, like a vacancy in Si, it can sample the configuration space much more efficiently than molecular dynamics [11]. However, it is not so exciting to be able to simulate a single vacancy hopping around in crystalline Si. Many methods could do as well or better than POP-ART for these simple cases.

The current challenge in computational physics is to develop an algorithm that can sample efficiently the phase space of complex and slow materials such as glasses, polymers and proteins. These materials typically suffer from a wide distribution of activation energy barriers and most numerical methods fail to move efficiently through the energy landscape of these systems.

A canonical system of this kind is amorphous silicon, a material with perfect four-fold coordination, yet disordered. Its energy landscape has been extensively studied using ART and similar methods [12–15]; one observation is that this material demonstrates a continuous distribution of activation energies.

We applied POP-ART to a model of a-Si at 600 K and find, as discussed in the next section, that the system spends most of its time moving between a small number of basins, resulting in an efficiency at best equivalent to MD, in terms of computational costs.

Because all steps in POP-ART are defined *on-the-fly*, it is possible to alleviate this problem efficiently and cheaply by introducing a memory kernel that keeps track of activated trajectories already visited and prevents visiting them again for a predetermined number of steps.

Before discussing the efficiency gain due to such a memory kernel, it is important to verify how much bias it introduces in the sampling. We do that on a simple model.

### 3.1. A simple thermodynamical model

The model we consider is as follows. On a square lattice with $L \times L$ sites and periodic boundary conditions, each of $N = L^2$ sites corresponds to a state. Assign to each state $i$, $i = 1, \ldots, N$, a (free) energy $E_i$ drawn from the uniform distribution in the range $[0, E_{\max}]$ (we use $E_{\max} = 1$ in our simulations). The probability that the system is located in state $i$ is then

$$P_i = \frac{1}{Z(\beta)} \exp(-\beta E_i), \tag{7}$$

where $\beta = 1/k_B T$, the inverse temperature. Here, $Z(\beta)$ is a normalization constant (the partition sum), set by the constraint $\sum_{i=1}^{N} P_i = 1$.

The dynamics of this model consists of jumps to one of the four nearest or four next-nearest neighbor sites. The transition rate $v_{ij}$ from state $i$ to state $j$ is given by

$$v_{ij} = v^*_{ij} \min(1, P_j/P_i), \tag{8}$$

where $v^*_{ij} = v^*_{ji}$. This model obeys detailed balance, i.e., each pair of states $i, j$ obeys by construction the relation

$$P_i v_{ij} = P_j v_{ji}. \tag{9}$$

To mimic typical situations encountered in POP-ART simulations of real complex systems, we chose the transition rates as follows. Of all transitions, 5% (chosen at random) are 'fast', with $v^*_{ij} = 1$, and the rest are 'slow', $v^*_{ij} = 10^{-3}$. Such a system will often be 'trapped', going back and forth between two states connected by a 'fast' transition. It should be possible to reduce this effect by avoiding moves that have occurred recently. This is the basis of an approach to global optimization (i.e., finding the global minimum of a given function) known as *tabu search* [16]. Using this approach to simulate thermodynamics of a system is potentially more problematic, as some systematic bias can be introduced – the probabilities of states will no longer be given by Eq. (7). However, if this bias is small, it might be worth tolerating it if significant speed-up is observed.

We study two possibilities of accelerating exploration of configuration space:

(1) The system is not allowed to return to any of the $H_s$ most recently visited states.
(2) The system is not allowed to repeat any of the $H_t$ most recent transitions.

We first study the question of systematic bias. We quantify it by

$$\Delta_b^2 = \frac{1}{N}\sum_{i=1}^{N}(\ln P_i - \ln f_i)^2, \tag{10}$$

where $f_i$ is the fraction of time during which state $i$ is observed. Generally speaking, the smaller $\Delta_b$, the smaller the error in the thermodynamic quantities that would be obtained in the simulation. Simulations proceed with continuous time Monte Carlo [17]. In the exact simulation (without memory), starting from state $i$ at time $t = 0$, first the time is incremented by

$$\Delta t_i = \left(\sum_j \nu_{ij}\right)^{-1}. \tag{11}$$

Next, at time $\Delta t_i$ one of the eight neighboring states is chosen; the probability that this state is $j$ is equal to $\Delta t_i \cdot \nu_{ij}$. Then the procedure is repeated for the new state $j$. In a simulation with memory, the procedure is similar, *except* the sum in Eq. (11) now runs over the allowed transitions only and likewise, the new state $j$ is chosen among the allowed states only.

We chose a system with $N = 100$ states and the memory length $H_s = H_t = 4$. (In fact, the results are not very sensitive to the choice of the memory length.) The results are shown in Fig. 2. For the first approach (memory of recently visited states), these results are obtained by running 500 simulations of 2000 steps each for a given realization of energies $E_i$ and rates $\nu_{ij}$ and then averaged over 500 such realizations. It is seen that $\Delta_b^2$ is about 1 in this case, which means that using this approach is out of the question – the bias is too strong. For the second approach (memory of recent transitions), the bias is much smaller and much longer simulations are needed to estimate it accurately. We have run 100 simulations of $3\times10^6$ steps or $5\times10^6$ steps each for every realization of energies and rates and then averaged over 100 realizations. Even simulations of this length are not sufficient for full convergence, and more accurate results were obtained by extrapolation. The bias $\Delta_b^2$ turns out to be on the order of $10^{-4}$, i.e., the observed probabilities of states are accurate to within about 1%, which should be acceptable in most cases.

To study the efficiency of sampling, we consider diffusion on the lattice. In place of time we need to choose a quantity that would be proportional to the CPU time in a POP-ART simulation. In POP-ART, the most time-consuming part is evaluation of the Jacobian for a transition. For this reason, the CPU time should be roughly proportional to the number of attempted transitions (both accepted and rejected), but not including those that are rejected because of the memory (as these will be rejected immediately, without evaluating the Jacobian). A convenient choice for the simulation procedure is then as follows: an allowed transition is selected with the probability proportional to $\nu^*_{ij}$ (rather than $\nu_{ij}$), but the selected transition is then accepted with the Metropolis probability $\min(1, P_j/P_i)$. The clock is then advanced by 1, regardless of whether the transition is accepted or rejected. Such a clock would indeed count all transition attempts, except those rejected because of the memory.

For the simulations, we have chosen a bigger $200\times200$ lattice to eliminate the effects its periodicity may have on the diffusion. A typical dependence of the mean squared

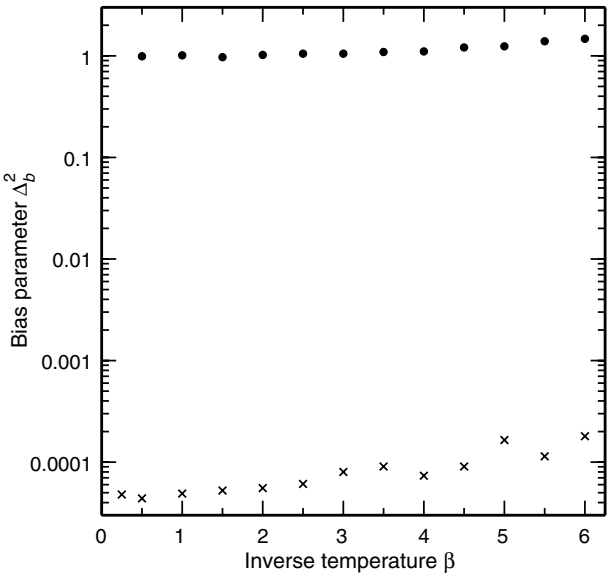


Fig. 2. The bias parameter $\Delta_b^2$ for the simple model considered in this paper for the approach with the memory of recently visited states (circles) and for the approach with the memory of recent transitions (crosses). All other parameters are given in the text.

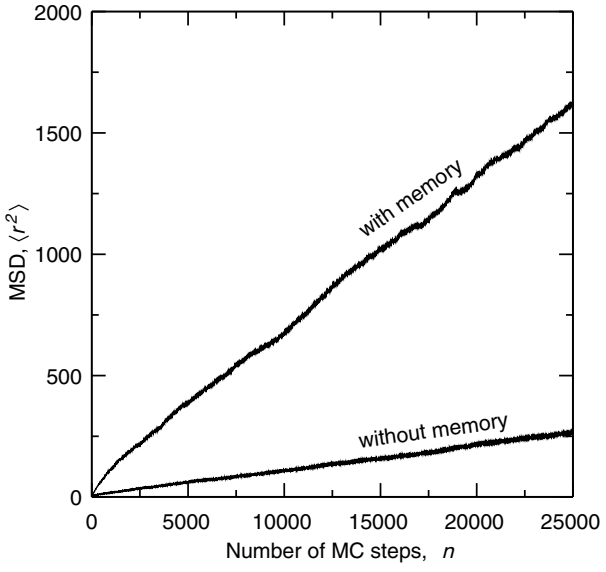


Fig. 3. The dependence of the mean squared displacement $\langle r^2\rangle$ on the Monte Carlo step number $n$ for the model discussed in the text at the inverse temperature $\beta = 5.0$, both without memory and with the memory of recent transitions. The data are averages over 1000 realizations with different sets of energies and rates.

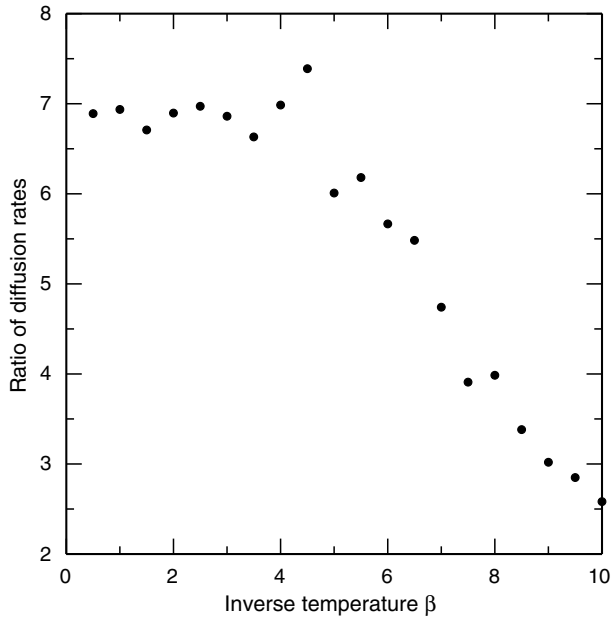


Fig. 4. The ratio of the diffusion constants obtained with the memory of recent transitions and without memory for the model described in the text. The diffusion constants are calculated as the ratios of the change in the mean squared displacement (averaged over 10 000 realizations), $\Delta\langle r^2\rangle$, over $\Delta n = 5000$ Monte Carlo steps, starting after 5000 steps for $\beta < 5$ and after 20 000 steps for $\beta \geqslant 5$.

displacement as a function of the number of Monte Carlo steps is shown in Fig. 3. After a short transition period, the motion is diffusive. It can be characterized by the diffusion constant, measured as the ratio $\Delta\langle r^2\rangle/\Delta n$, where $r$ is the distance from the state occupied at time $t = 0$ (with the lattice constant equal to 1), $n$ is the Monte Carlo step number, and the average is over different realizations. Fig. 4 shows the ratio of the diffusion constants obtained with and without memory. It is seen that for the particular set of parameters chosen and not too low temperatures, there is roughly a sevenfold improvement in the diffusion rate, although the ratio decreases with decreasing temperature.

## 4. Application to POP-ART

As mentioned above, we test POP-ART with and without the memory kernel using a 1000-atom model of well-relaxed amorphous silicon, described using a modified Stillinger–Weber potential [18] and run at 600 K, well below melting (about 1700 K). For comparison we also run a constant-temperature molecular dynamics simulation using a time step of 1 fs. In the POP-ART simulations, we use as threshold for the eigenvalue $\lambda_0 = -1.5\ \text{eV/Å}^2$.

To establish the efficiency of the method, we measure its ease of sampling the phase space. To compare the different methods we periodically relax the 600-K configurations into their local minimum, dubbed *inherent structure* [19]. For POP-ART simulations, the configurations are relaxed after every activated event; for MD, at intervals of 50 ps.

Results for the standard POP-ART are presented in Fig. 5. In these graphs, we compare two independent POP-ART runs with the MD simulation. We see that, as a function of the number of force evaluations, which is directly proportional to the real computational cost, POP-ART is, at best, as efficient at molecular dynamics if we consider only the RMSD. It does better than or as well as MD over $2 \times 10^8$ force evaluations with respect to the inherent structure energy, however. This is already positive as it includes the additional cost of computing the activated path and the Jacobian.

It is possible to improve significantly the efficiency of POP-ART by adding the memory kernel discussed above. We use the version of the memory kernel with the transition history, rather than state history. Besides being much less biased, as shown for the simple model in the previous section, this approach is also convenient for other reasons. First, while it is easy to find out if a given inherent structure has already been visited, it is not so easy for basins, as a basin may contain several inherent structures. Second, a particular transition can be identified repeatedly even if many events happen elsewhere in the system between the instances of this transition. This is an important point,

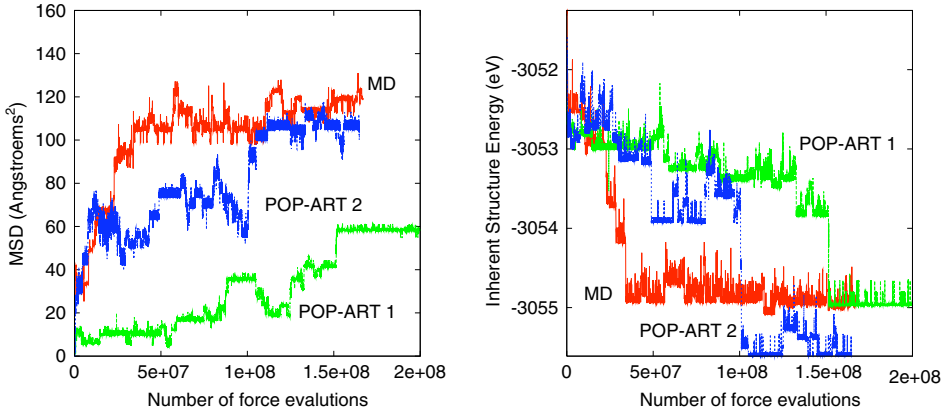


Fig. 5. The total mean squared displacement (MSD) and the potential energy of the inherent structure configurations as a function of the total number of force evaluations for a standard molecular dynamics simulation and two standard POP-ART simulations of a 1000-atom amorphous silicon model at 600 K. The configurations are minimized every 50 ps or after every event, respectively.

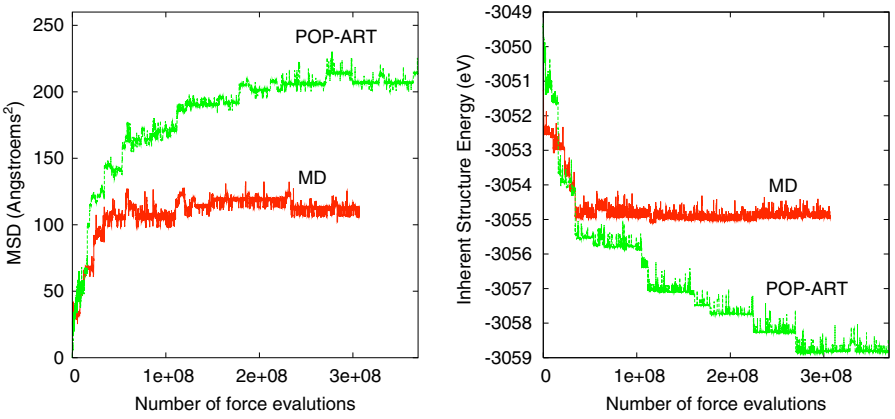


Fig. 6. The total mean squared displacement (MSD) and the potential energy of the inherent structure configurations as a function of the total number of force evaluations for a standard molecular dynamics simulation and a POP-ART simulation with a memory kernel of 250 eigenvectors and an overlap threshold of 60% of a 1000-atom amorphous silicon model at 600 K. The configurations are minimized every 50 ps or after every event, respectively.

since, particularly for condensed matter systems, individual events are usually localized in a relatively small part of the system and the Hessian eigenvectors corresponding to these events consequently span only a small part of the system's coordinates. This ability to identify events using local information, ignoring what occurs in other parts of the system, is a unique feature of POP-ART.

Fig. 6 shows the MSD and total energy for the inherent structures for a POP-ART simulation with the same threshold as above, but with a memory kernel of 250 eigenvectors. While the size of the memory kernel is much bigger than that used in the simple model of the previous section, one should keep in mind that the total number of transitions out of a given basin is $O(N)$, so this size is still a relatively small fraction of all transitions out of a basin. For the purpose of this simulation, we define two eigenvectors as identical if the $3N$-dimensional vectors overlap by at least 60%. We see that while molecular dynamics reaches a plateau in the MSD and the energy after about 50 ns (or $5\times10^7$ force evaluations), POP-ART with memory continues to explore the phase space without reaching a plateau. This is particularly clear on the energy graph.

## 5. Conclusion

This paper presents an extension to the properly obeying probability activation-relaxation technique (POP-ART), which uses the local (without reference to the local energy-minimum) definition of moves in this method to implement an efficient memory kernel. We can show, using a simple model, that the use of this kernel provides a more efficient sampling of the phase space at the cost of some bias, which, however, in the case of memory of recent transitions is shown to be rather small.

Using a model of amorphous silicon, we have shown that the memory kernel can help significantly in exploring the phase space of this disordered material at low temperature. While standard POP-ART is as efficient as MD, the use of the memory kernel increases its efficiency significantly. POP-ART with the memory kernel is a promising method for disordered materials, systems for which very few efficient methods exist. We plan to apply this method to a number of such systems.

## Acknowledgements

M.C. acknowledges partial support from the Ministère de l'éducation, du sport et des loisirs (Québec). N.M. would like to thank the Institute for Theoretical Physics at Utrecht University for its hospitality and the Dutch NWO for financial support. N.M. is supported in part by FQRNT (Québec), NSERC and the Canada Research Chair program.

## References

[1] E. Marinari, G. Parisi, Europhys. Lett. 19 (1992) 451.
[2] J. Lee, Phys. Rev. Lett. 71 (1993) 211.
[3] F. Wang, D. Landau, Phys. Rev. Lett. 86 (2001) 2050.
[4] A.F. Voter, Phys. Rev. Lett. 78 (1997) 3908.
[5] X. Wu, S. Wang, J. Chem. Phys. 110 (1999) 9401.
[6] G.T. Barkema, N. Mousseau, Phys. Rev. Lett. 77 (1996) 4358.
[7] J.P.K. Doye, D.J. Wales, Z. Phys. D 40 (1997) 194.
[8] L.J. Munro, D.J. Wales, Phys. Rev. B 59 (1999) 3969.
[9] G. Henkelman, H. Jónsson, J. Chem. Phys. 111 (1999) 7010.
[10] R. Malek, N. Mousseau, Phys. Rev. E 62 (2000) 7723.
[11] H. Vocks, M.V. Chubynsky, G.T. Barkema, N. Mousseau, J. Chem. Phys. 123 (2005) 244707.
[12] N. Mousseau, G.T. Barkema, Phys. Rev. E 57 (1998) 2419.
[13] G.T. Barkema, N. Mousseau, Phys. Rev. B 62 (2000) 4985.
[14] Y. Song, R. Malek, N. Mousseau, Phys. Rev. B 62 (2000) 15680.
[15] F. Valiquette, N. Mousseau, Phys. Rev. B 68 (2003) 125209.
[16] F. Glover, M. Laguna, Tabu Search, Kluwer Academic, Dordrecht, 1997.
[17] M.E.J. Newman, G.T. Barkema, Monte Carlo Methods in Statistical Physics, Oxford University, Oxford, 1999.
[18] R.L.C. Vink, G.T. Barkema, W.F. van der Weg, N. Mousseau, J. Non-Cryst. Solids 282 (2001) 248.
[19] F.H. Stillinger, T.A. Weber, Science 225 (1984) 983.